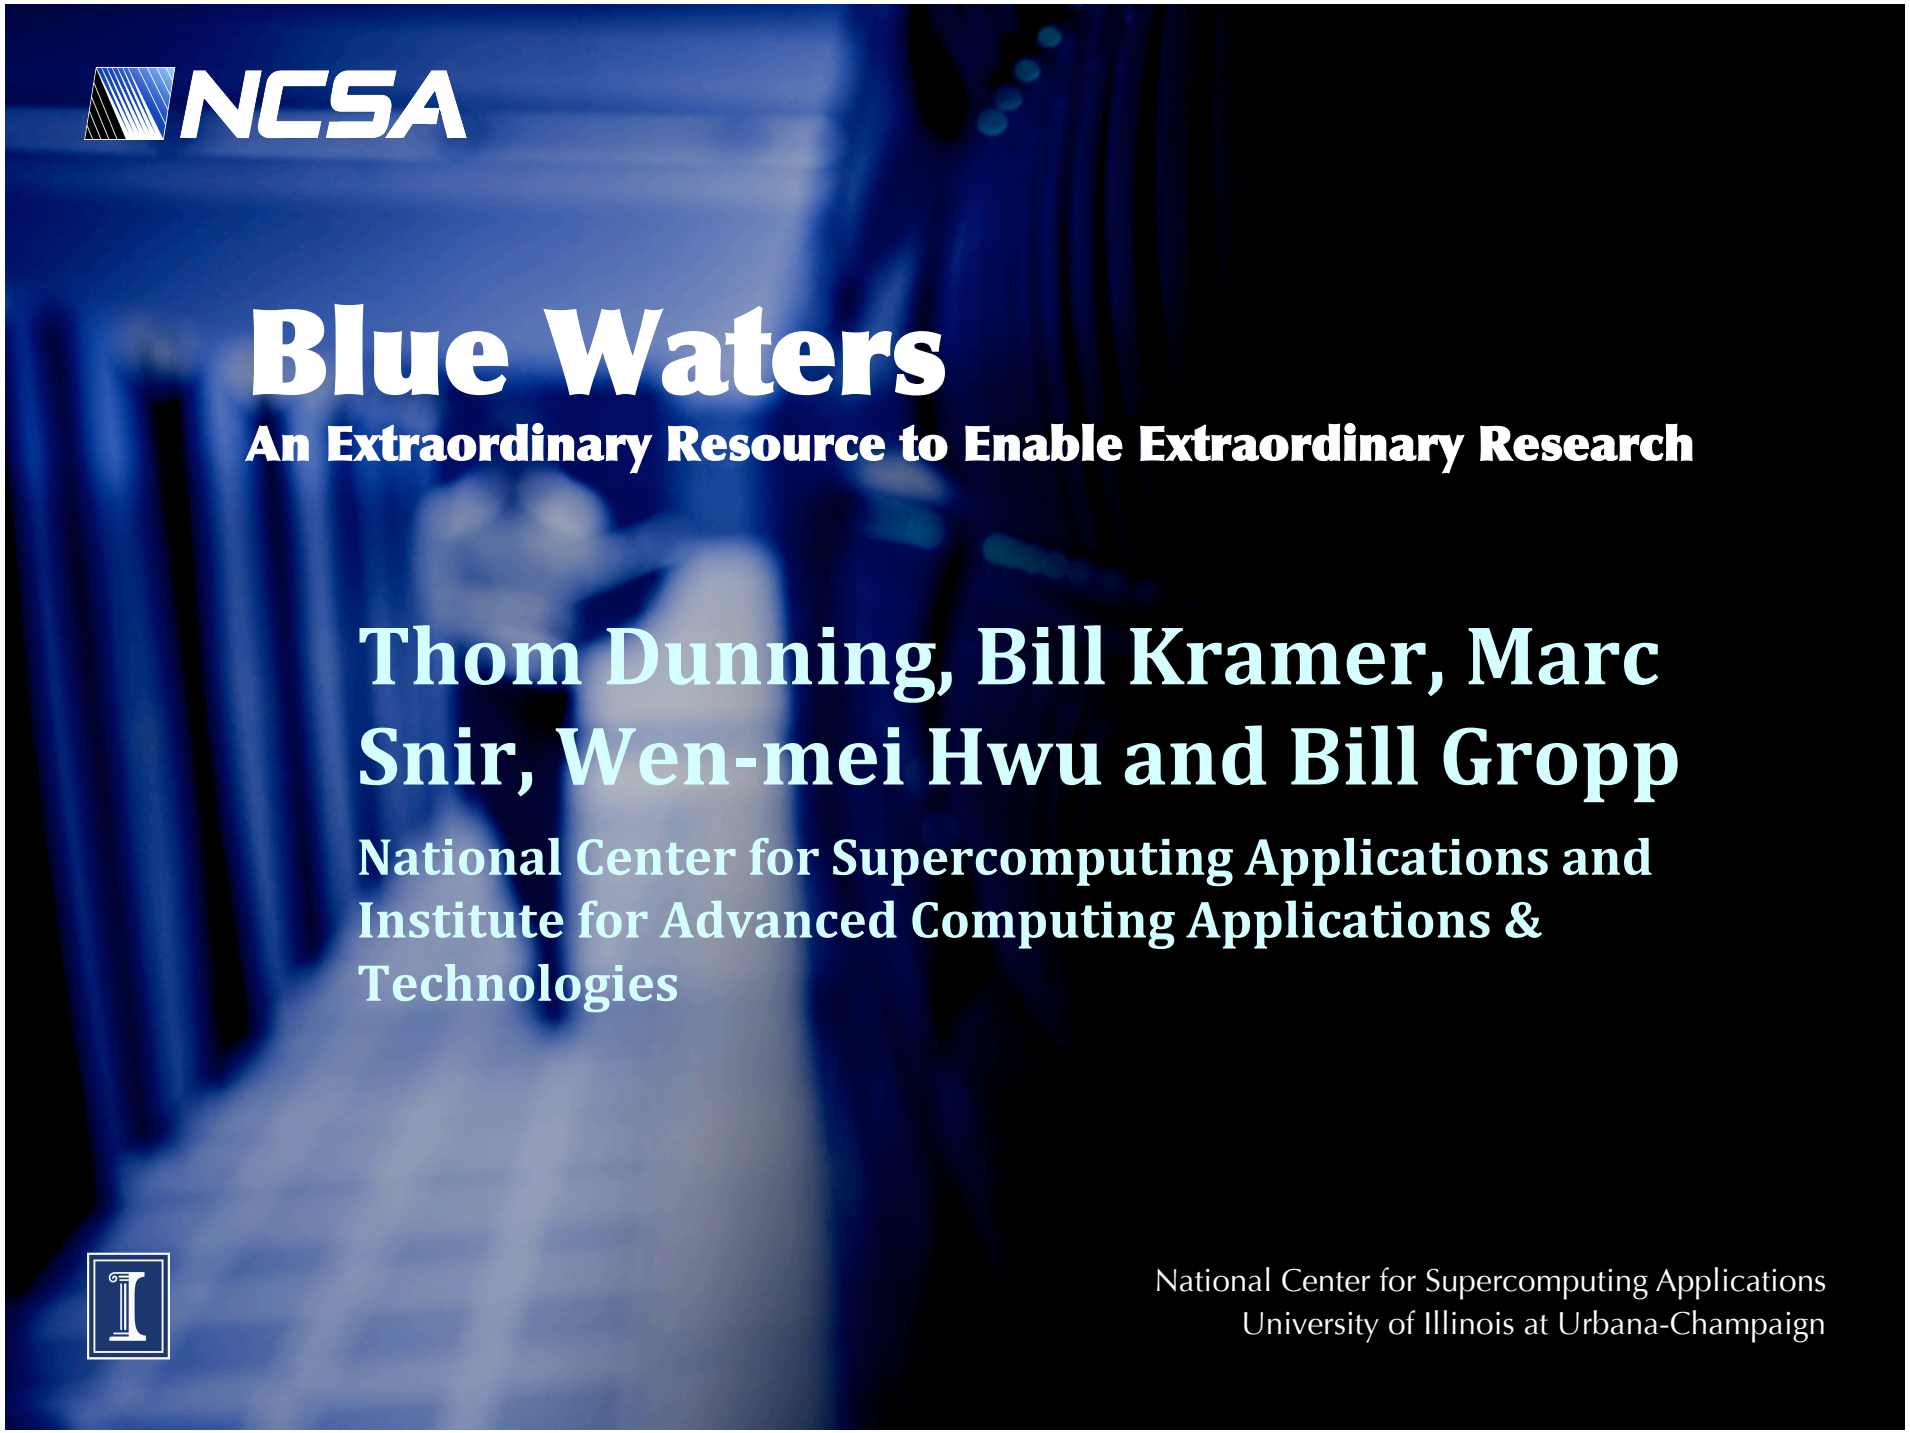

Blue Waters

An Extraordinary Resource to Enable Extraordinary Research

Thom Dunning, Bill Kramer, Marc Snir, Wen-mei Hwu and Bill Gropp

National Center for Supercomputing Applications and Institute for Advanced Computing Applications & Technologies

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

Computational Science and Engineering

Increased computing power enables advances in a broad range of science and engineering disciplines:

Molecular Science

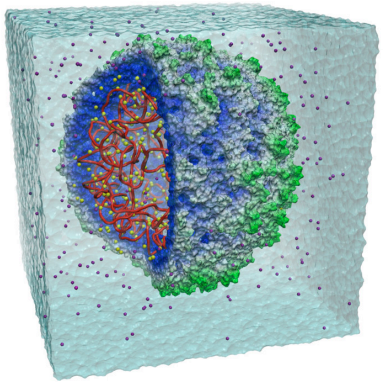

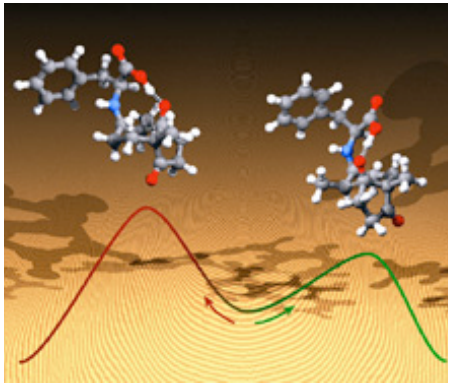

Weather & Climate Forecasting

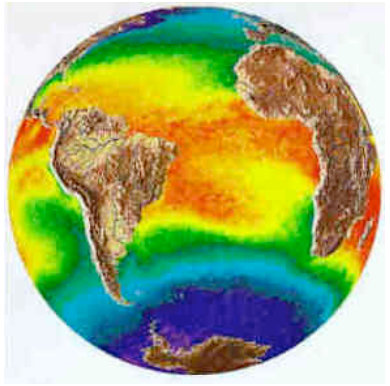

Astronomy

Earth Science

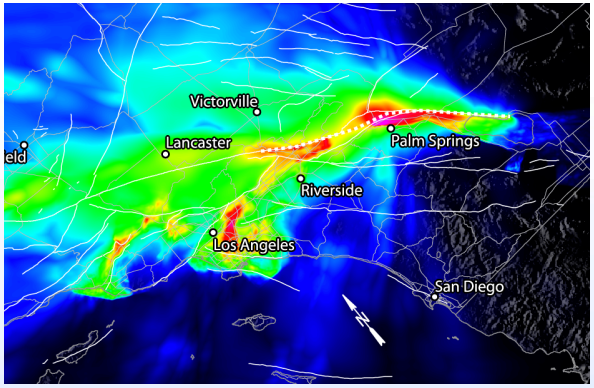

Health

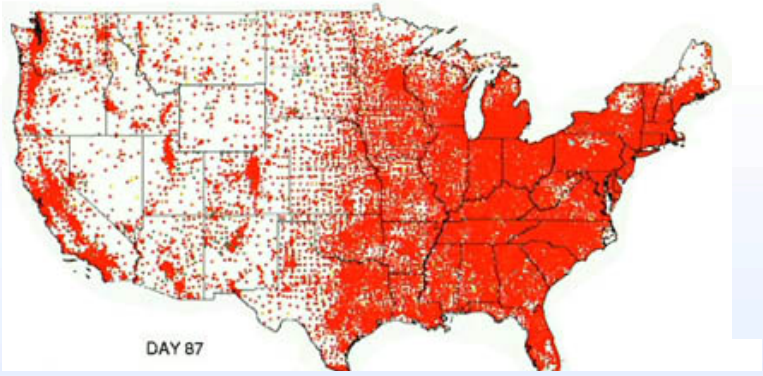

NSF High-end Computing Strategy

- **Three Resource Levels**
  - Track 3: University owned and operated
  - Track 2: Several NSF-funded supercomputer & specialized computing centers (TeraGrid)
  - Track 1: NSF-funded leading-edge computer center
- **Computing Resources**
  - Track 3: 10s–100s TF
  - Track 2: 500–1,000 TF, ~100+ TB of memory
  - Track 1: *see following slides*

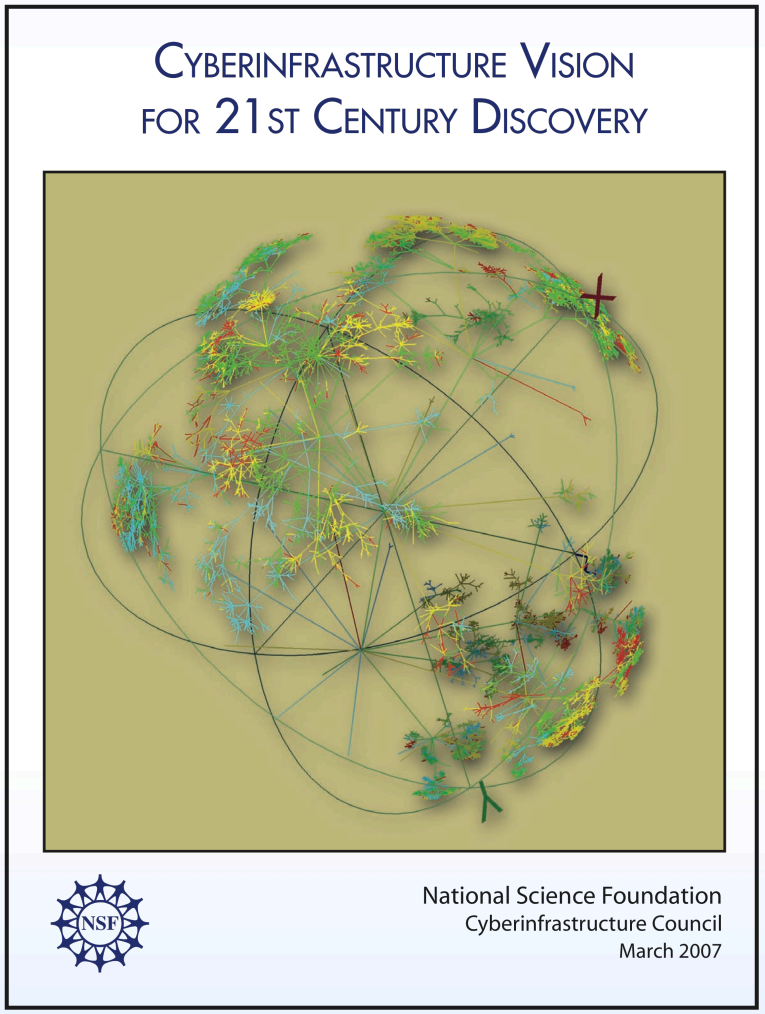

Blue Waters Project

Goals of Blue Waters Project

- **Achieve Sustained Petascale Performance**
  - Petascale performance on broad range of science & engineering applications
  - Acceptance criteria include sustained performance on petascale applications and applications that use only fraction of system
    - Tests include time to read data and write the results
- **Provide Clear Pathway to Petascale Computation**
  - Robust, powerful petascale computing system
    - *See following slides*
  - Robust, powerful petascale computing software
    - Existing HPC stack
    - Adaptation of existing tools, techniques, libraries, etc. to take full advantage of petascale computers; exploration of new tools, programming models, etc. to facilitate performance, software development, operation of system, etc.
  - Collaboration with Approved Research Teams to optimize their applications

IBM Programmable Easy-to-use Reliable Computing System

November 24, 20006

IBM wins DARPA funding

HPCS program allows IBM to pursue a vision of petascale computing systems

IBM's Programmable Easy-to-use Reliable Computing System (PERCS) was selected by the Defense Advanced Research Projects Agency (DARPA) as one of two system designs to be developed and demonstrated as part of phase III of the High Productivity Computing Systems program (HPCS). *Such designs must support the eventual scaling of sustained computation to 10 petaflops and a software environment that enables domain experts to effectively use that computing power*. These capabilities are required by HPCS to meet the need for commercially successful petascale computing systems for high-end users in government, science and industry in 2010. …

PERCS will meet these goals with a scalable system based on future POWER series technologies. The PERCS program will substantially increase the research and development activities in IBM technologies planned for 2010 and beyond. These will enable IBM to meet the HPCS goals and enhance the capabilities of IBM's line of business systems. This will entail IBM making significant investments in the next generation of [a number of] technologies.

…

Attributes of Blue Waters

- **Compute-intensive Performance**
  - … maximum core performance to minimize number of cores needed, lessen the impact of sections of code with limited scalability
  - … low latency, high-bandwidth communications fabric to facilitate scaling to large numbers of compute cores
- **Memory-intensive Performance**
  - … large amount of memory
  - … low latency, high-bandwidth memory subsystem to enable rapid transfer of data from the cores to memory
- **Data-intensive Performance**
  - … large quantity of on-line disk, massive quantity of archival storage
  - … high-bandwidth I/O subsystem to enable rapid transfer of data from memory to on-line disks and then to archive
- **Reliable Operation**
  - … mainframe RAS technologies

Some Unique Features of Blue Waters

- **Shared/Distributed Memory Computing System**
  - Powerful Shared Memory Multichip Module (MCM)
  - New high bandwidth fabric interconnects all nodes
  - Hardware support for global shared memory
- **High Performance I/O Subsystem**
  - More than an order of magnitude faster than existing I/O subsystems
  - On-line disks fully integrated with archival storage system
- **Broad Range of Systems**
  - IBM Power7 systems will range from servers to supercomputers to address science and engineering problems at all levels

Power 7 Chip: Heart of Blue Waters

- **Base Technology**
  - 45 nm, 576 mm^2
  - 1.2 B transistors
- **Chip**
  - 8 cores
  - 12 execution units/core;
    up to 256 GF peak performance
  - 4-way SMT/core
  - Caches
    - 32 KB I-, D-cache, 256 KB L2/core
    - 32 MB L3 (private/shared)
  - Dual DDR3 memory controllers
    - 128 GB/s memory bandwidth
- **Multichip Module**
  - Four Power7 chips plus Hub chip
  - 128 Gbytes of memory
  - 192 GB/s bandwidth off-chip

Power7 Chip

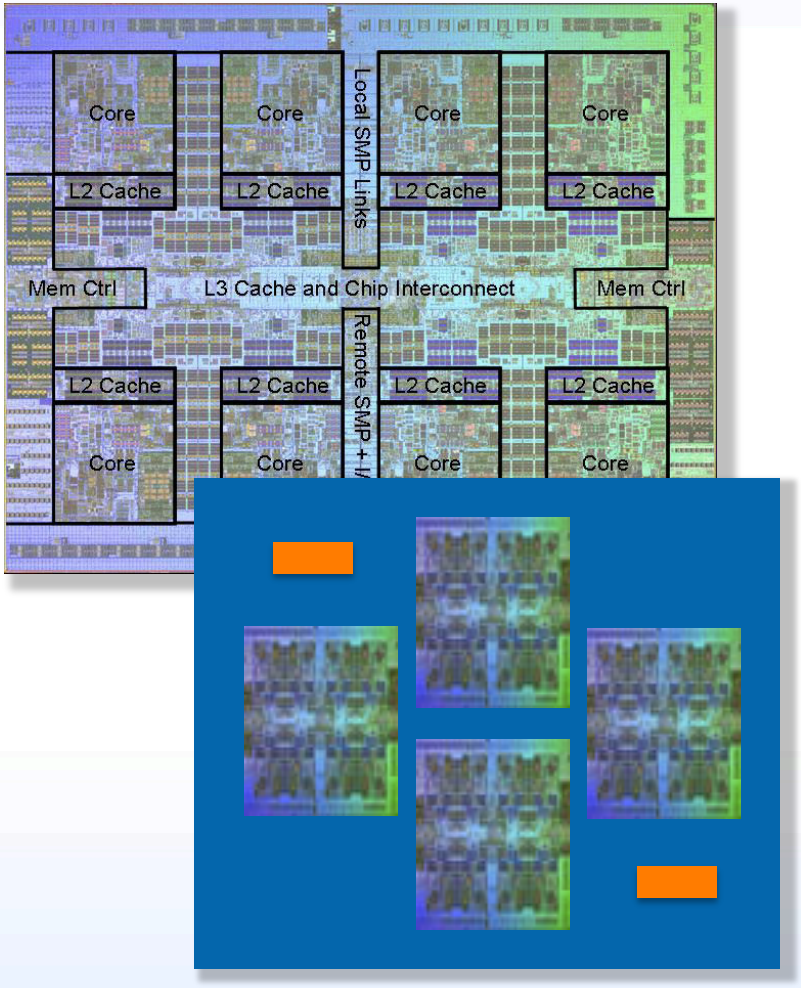

Quad-chip MCM

BW Drawer: Basic Compute Unit

PCIe Cards | Hub Chips | Memory | MCMs | Memory | Power Supply

Drawer: 2U
39w x 72d inches
Completely water cooled (MCM, Hub, memory)

Building Blue Waters

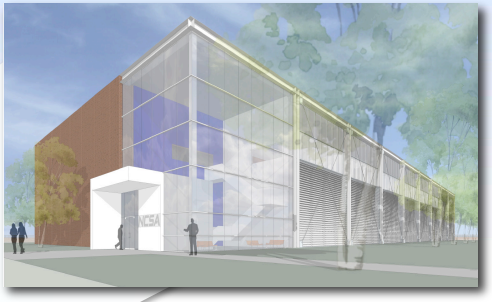

Blue Waters
~ 1 PF sustained
> 200,000 cores
> 800 TB of memory
> 10 PB of disk storage
> 500 PB of archival storage
> 100 Gbps connectivity (Chicago)

Blue Waters Building Block
32 IH drawers
32 TB memory
256 TF (peak)
4 Storage systems
10 Tape drive connections

IH Drawer
8 QCM's (256 cores)
8 Hubs
1 TB memory
8 TF (peak)

Fully water cooled

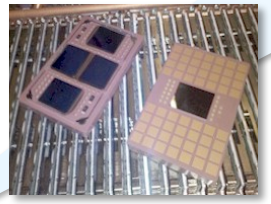

Multi-chip Module
4 Power7 chips
128 GB memory
512 GB/s memory bandwidth
1 TF (peak)

Hub Chip
1,128 GB/s bandwidth

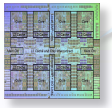

Power7 Chip
8 cores, 32 threads
L1, L2, L3 cache (32 MB)
Up to 256 GF (peak)
45 nm technology

Petascale Computing Facility

Partners

EYP MCF/
Gensler
IBM
Yahoo!

- **Modern Data Center**
  - 90,000+ ft^2 (total)
  - 20,000 ft^2 machine room (net)
- **Energy Efficiency**
  - LEED certified (gold)
  - PUE < 1.2

www.ncsa.uiuc.edu/BlueWaters

Petascale Resource Computing Allocations

Approved Science & Engineering Research Teams

Petascale Computing Resource Allocations

- **Solicitation: NSF 08-529 (PRAC)**
- **Selection Criteria**
  - Compelling science or engineering research question
  - Question that can only be answered using a system of the scale of Blue Waters (cycles, memory, I/O bandwidth, etc.)
  - Evidence, or a convincing argument, that the application code can make effective use of Blue Waters
  - Source (or sources) of funds to support the research work and any needed code development effort
- **Funding Level**
  - Allocation or provisional allocation of time on Blue Waters
  - Travel funds to enable teams to work closely with Blue Waters Project team at NCSA
- **Next Due Date**
  - March 17, 2010 (*annually thereafter*)

Engagement with Research Teams

- **Provide Details on Blue Waters System**
- **Provide Assistance with Blue Waters Software**
  - Numerical libraries
  - MPI, OpenMP, ARMCI/Global Arrays, LAPI, OpenSHMEM, Charm++
  - Compilers (Fortran, C, UPC, Co-array Fortran)
- **Provide Assistance with Blue Waters Hardware**
  - Chip and network simulators
  - Staged access to Power7 hardware
- **Provide Training**
  - On-line documentation
  - Webinars/tutorials/on-line courses (~8 per year)
  - Workshops (~2 per year)

PRAC Awards to Date

- **Biological Sciences**
  1. Computational Microscope
     - Klaus Schulten, Laxmikant Kale, University of Illinois at Urbana-Champaign
  2. Petascale Simulations of Complex Biological Behavior in Fluctuating Environments
     - Ilias Tagkopoulos, University of California, Davis
- **Engineering**
  3. Petascale Computations for Complex Turbulent Flows
     - Pui-Kuen Yeung, James Riley, Robert Moser, Amitava Majumdar, Georgia Institute of Technology

PRAC Awards to Date *(cont'd)*

- **Geosciences**

4. Petascale Research in Earthquake System Science on Blue Waters
   - Thomas Jordan, Jacobo Bielak, University of Southern California
5. Testing Hypotheses about Climate Prediction at Unprecedented Resolutions on the Blue Waters System
   - David Randall, Ross Heikes, Colorado State University; William Large, Richard Loft, John Dennis, Mariana Vertenstein, National Center for Atmospheric Research; Cristiana Stan, James Kinter, Institute for Global Environment and Society; Benjamin Kirtman, University of Miami
6. Understanding Tornados and Their Parent Supercells Through Ultra-High Resolution Simulation/Analysis
   - Robert Wilhelmson, Brian Jewett, Matthew Gilmore, University of Illinois at Urbana-Champaign
7. Enabling Large-Scale, High-Resolution, and Real-Time Earthquake Simulations on Petascale Parallel Computers
   - Liqiang Wang and Po Chen, University of Wyoming

PRAC Awards to Date *(cont'd)*

- **Mathematics & Physical Sciences**
  - Astronomical Sciences
    8. Computational Relativity and Gravitation at Petascale: Simulating and Visualizing Astrophysically Realistic Compact Binaries
       - Manuela Campanelli, Carlos Lousto, Hans-Peter Bischof, Joshua Faber, Yosef Ziochower, Rochester Institute of Technology
    9. Enabling Science at the Petascale: From Binary Systems and Stellar Core Collapse to Gamma-Ray Bursts
       - Eric Schnetter, Gabrielle Allen, Mayank Tyagi, Peter Diener, Christian Ott, Louisiana State University
    10. Formation of the First Galaxies: Predictions for the Next Generation of Observatories
       - Brian O'Shea, Michigan State University; Michael Norman, University of California at San Diego

PRAC Awards to Date (*cont'd*)

- **Mathematics & Physical Sciences**
  - Astronomical Sciences
    11. Peta-Cosmology: Galaxy Formation and Virtual Astronomy
        - Kentaro Nagamine, University of Nevada at Las Vegas; Jeremiah Ostriker, Princeton University; Renyue Cen, Greg Bryan
    12. Petascale Simulation of Turbulent Stellar Hydrodynamics
        - Paul Woodward, Pen-Chung Yew, University of Minnesota, Twin Cities
  - Chemistry
    13. Computational Chemistry at the Petascale
        - Monica Lamm, Mark Gordon, Theresa Windus, Masha Sosonkina, Brett Bode, Iowa State University
    14. Super Instruction Architecture for Petascale Computing
        - Rodney Bartlett, Erik Duemens, Beverly Sanders, University of Florida; Ponnuswamy Sadayappan, Ohio State University

PRAC Awards to Date *(cont'd)*

- **Mathematics & Physical Science**
  - Materials Research
    15. Breakthrough Petascale Quantum Monte Carlo Calculations
        - Shiwei Zhang, College of William and Mary
    16. Electronic Properties of Strongly Correlated Systems Using Petascale Computing
        - Sergey Savrasov, University of California, Davis; Kristjan Haule, Gabriel Kotliar, Rutgers University
  - Physics
    17. Lattice QCD on Blue Waters
        - Robert Sugar, University of California at Santa Barbara

PRAC Awards to Date *(cont'd)*

- **Social, Behavioral and Economic Sciences**
  18. Simulation of Contagion on Very Large Social Networks with Blue Waters
      - Keith Bisset, Xizhou Feng, Virginia Polytechnic Institute and State University

Questions?